Reducing System Power With A New NAND Flash Memory Interface

by Peter Gillingham, Vice President and Chief Technology Officer
MOSAID Technologies Inc

Power consumption in NAND Flash memory subsystems, particularly those with large numbers of memory devices such as SSDs (Solid State Drives), is typically dominated by the I/O power dissipated in transferring data between the controller and the memory devices. Achieving low power is an important consideration. In mobile devices, such as cell phones and media players, extended battery life or smaller form factors can be achieved with reduced power mass storage devices. In enterprise applications such as SSD storage arrays, low power NAND devices provide a reduced thermal budget which leads to higher computational density and reduced energy costs.

A new ring topology that delivers significant power benefits over the traditional bus topology can be employed by all mainstream memory devices, including NAND Flash, NOR Flash, SRAM and DRAM. An SSD implemented with this architecture can use half the power of an equivalent SSD employing an ONFI (Open NAND Flash Interface), 2.0 NAND Flash device. This is significant in both portable applications where battery life is critical and in enterprise applications where heat and energy footprint considerations dominate.

The Challenge

NAND Flash products were first introduced by Toshiba in the early 1990s. The interface has remained virtually unchanged to the present day and has been accepted by the industry as a de-facto standard. In 2006 ONFI was formed by manufacturers and users of NAND Flash memory to resolve incompatibilities in legacy NAND Flash and to enhance the performance of future NAND Flash devices. In early 2008, ONFI released specification 2.0 which described a synchronous DDR NAND Flash device supporting clock frequencies up to 66 MHz and providing peak bandwidth of 133 Mbit/s.

ONFI 2.0 supports up to 16 loads. Any device, whether a memory chip or the controller, which wants to send data over the bus to another device, must drive all 16 other devices connected to the bus. This wastes power. To drive all these loads at the 133 Mbit/s data-rate a powerful 18 Ω output driver is required. This driver is roughly three times more powerful than the output driver in legacy NAND products. A larger transistor has a larger parasitic capacitance which represents additional loading on the bus. Pushing a multi-drop bus structure to its limits in this manner leads to a power-performance death spiral. For higher performance the output driver must be enlarged, resulting in higher pin capacitance (and higher power), necessitating further increases in output driver size.

A New Topology

A new interface approach for mass storage applications is a unidirectional ring in which data are passed from device to device over point-to-point connections. Although each device adds one clock cycle of latency, this is not an issue in mass storage applications. The latency around the ring is typically less than 1% of the time it takes to read or write the Flash memory device. The first generation topology employs a common DDR clock distributed to all memory devices and achieves performance up to 266 Mbit/s over an 8-bit bus.

On power up, each device in the ring is assigned a unique device identifier, starting with 0 and incrementing by 1 up to a maximum of 254. In operation, commands including the target device number originate from the controller and pass through a number of devices until they reach the addressed device. Write data also originating from the controller are appended to the write command packet and truncated by the addressed device. Unlike the multi-drop bus topology, data need not be sent any further than the intended destination. This reduces power because the pin capacitances of the devices further down the ring from the addressed device do not see any data transitions.

Similarly, with read data, only the portion of the ring from the addressed device back to the controller is active. If reads and writes are evenly distributed among the devices, on average half of the bus would be dissipating power. This represents a fundamental power advantage of the ring topology over multi-drop bus topologies. In mass storage applications such as SSD, read operations typically occur more frequently than write operations. A system designer can further optimize system power by storing the most frequently read information in the last device in the ring.

A further advantage of the ring topology over the bus topology is the ability to have simultaneous read and write data transactions. If the device being written to is located upstream of the device being read, there is no bus contention. A sophisticated controller can schedule read and write transactions to maximize the time during which simultaneous operations can occur, and thereby increase the aggregate ring bandwidth.

Power Comparison

Table 1 summarizes the key parameters of ONFI 2.0 and ring topology that are required to calculate I/O power. Four packages, each including four die, are clocked at the maximum 66 MHz clock rate to achieve 133 Mbit/s bandwidth. For apples-to-apples comparison, we use the same number of packages and die for all configurations, even though many more loads could be accommodated by the new ring topology.

	NAND	ONFI 2.0	Ring Topology	units
Data rate	40	133	266	Mb/s/pin
# signal pins	15	16	24	
I/O voltage	3.3	1.8	1.8	V
Output drive strength	50	18	50	Ω
Output capacitance	4	5	3	pF
Input capacitance	8	5	2	pF
# devices	4	16	255	

Table 1: Legacy NAND, ONFI 2.0, And Ring Topology Parameters

In calculating CV power for the two configurations, we assume that the bus is occupied 100% of the time with read and write data transactions. In reality, some time is required for commands and gaps between data bursts, but this is negligible in most NAND Flash applications which are characterized by relatively long data bursts. In addition we assume random data with a 50% probability of a data bit transition in any one bit period.

For ring topology we assume a uniform probability of addressing any device, so that on average the data packet will travel half way around the ring. Clock power is also included in these calculations.

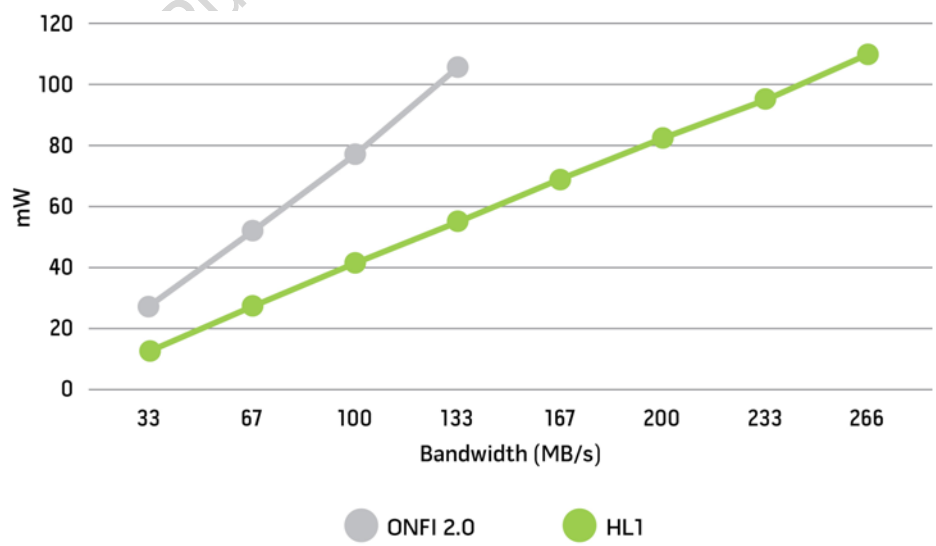

Fig. 1: Ring Technology Vs ONFI 2.0 I/O Power As Function Of Data Rate

Fig. 1 shows I/O power consumed by the two topologies with a fixed number of memory devices. The new architecture can operate up to 266 Mbit/s, twice the maximum data-rate as ONFI 2.0, but at 133 Mbit/s it consumes about one-half the power. This is a significant advantage considering that core power is a fraction of I/O power at these data-rates.

Re-Thinking NAND Flash Architecture

It's time to re-think the architecture for mass storage devices such as NAND Flash. Legacy NAND Flash adopted the multi-drop bus architecture employed by DRAM and SRAM. The multi-drop topology is appropriate for main memory and code storage applications where latency is most important, even at the cost of additional power. In mass storage applications characterized by large data transfers, a few additional clock cycles of latency have absolutely no effect on performance. But next generation NAND Flash will break away from the traditional multi-drop bus architectures and embrace the ring topology to deliver reduced power and higher throughput.

About The Author

Peter Gillingham is the Vice President and Chief Technology Officer of MOSAID Technologies Inc and one of the architects of HL-NAND, a new standard for high-performance Flash memory. A prolific inventor with over 60 issued US patents, Peter can be reached at gillingham@mosaid.com To learn more about the new HL-NAND technology architecture, visit http://www.hlnand.com